Architecture and Technology Program

March 2003

Offline Archive Media Trade Study

USGS
science for a changing world

Offline Archive Media Trade Study

March 4, 2003

By SAIC

Prepared by:

Tom Bodoh Date
Principal Systems Engineer
SAIC

Approved by:

Tony Butzer Date
Architecture & Technology Task Lead
SAIC

Reviewed by:

Ken Gacke Date
Principal Systems Engineer
SAIC

Reviewed by:

Cheryl Greenhagen Date
Principal Systems Engineer
SAIC

Reviewed by:

Al Engelbrecht Date
Principal Electronic Engineer
SGT

Reviewed by:

Kenneth Boettcher Date
Archive Manager
SAIC

Approved by:

John Faundeen Date
Archivist
USGS

Preface

This document contains the Offline Archive Media Trade Study for the Archive task. The Trade Study presents the background, technical assessment, test results, and the follow up recommendations as required by the Architecture and Technology Task Lead.

This Trade Study was prepared by

SAIC (Science Applications International Corporation)
USGS/EROS Data Center
Sioux Falls, SD 57198

Team members:

Tom Bodoh
Ken Boettcher
Ken Gacke
Cheryl Greenhagen
Al Engelbrecht
Tony Butzer

Abstract

This document is a trade study comparing offline digital archive storage technologies. The document compares and assesses several technologies and recommends which should be deployed as the next generation standard for the USGS at the EROS Data Center (EDC). Archives must regularly evolve to the next generation of digital archive technology and the technology chosen must remain reliable until the next migration. Note that this study is a revisit of a study completed in FY01 (Fiscal Year 2001).

Contents

1.0 Introduction

1.1 Purpose and Scope

This document provides an assessment of the options for the next generation of offline digital archive storage technology to be used for the Digital Archives of the USGS. The selected technology must be capable of safely retaining data until space, cost, and performance considerations would drive migration.

It is envisioned that within five years, most or all of the USGS archive holdings will reside on nearline storage and will be backed by an offline master copy and an offsite copy. The nearline copy is referred to as the working copy.

1.2 Background

The USGS, Earth Resources Observation Systems (EROS) Data Center, located in Sioux Falls, SD, currently archives offline datasets using several technologies. In 1992, the TMACS (TM/MSS Archive Conversion System) system was deployed to transcribe Landsat archives from HDT (High Density Tape) to DCT (Digital Cassette Tape). Both HDT and DCT utilize large, expensive analog instrumentation drives, which require frame synchronization, driving the cost of transcribing Landsat HDTs to DCTs to exceed $1,000,000 for each generation of media. Note that DCT and HDT are not purely analog. Although the crucial IRIG (InteRange Instrumentation Group) data is stored in analog format, the image data is stored in digital format. Though much of the conversion from HDT has been completed, additional HDT tapes were recently received. All HDT tapes transcribed to DCT by TMACS have been retained since no backup copies of the DCT tapes have been made.

Locating, rehabilitating, and integrating HDT drives has been costly in terms of labor, parts, and vendor service costs. The ongoing maintenance costs for the HDT and DCT drives are excessive since there is little industry experience and only a single vendor to support each brand of drive. The HDT and DCT drives in existence today number in the dozens, with the count decreasing each year as other users transition to digital media.

The "technology of choice" for the USGS archives has been the 35 GB (Gigabyte) DLT (Digital Linear Tape) 7000 for the past five years. Two new transcription systems were implemented in the past four years, transcribing HDT media to computer compatible DLT 7000. The DLT 7000 drive was retired by Quantum two years ago, replaced by the DLT 8000. The DLT 8000 has not been widely accepted since the SuperDLT drives had already been pre-announced when the DLT 8000 was released. A USGS study of DLT 7000 errors revealed that they exhibit a greater percentage of data loss as compared to 3480, 3490 and 9840.

Table 1-1 summarizes the offline archive tape technologies currently in use:

Tape Drive Technology	Capacity	Transfer rate	Type
HDT	3.4 GB	10.6 MB/sec	Analog
3480	200 MB	2 MB/sec	Digital
3490	900 MB	2.7 MB/sec	Digital
DLT 7000	32 GB	4.7 MB/sec	Digital
DCT (Ampex DCRsI)	45 GB	12 MB/sec	Analog
SuperDLT 220	98.8 GB	8.1 MB/sec	Digital

Table 1-1 Past and current archive technologies used

The USGS has utilized SuperDLT 220 extensively for onsite and offsite backups, and the LP-DAAC (Land Processes Distributed Active Archive Center) has used it to archive MODIS (Moderate Resolution Imaging Spectro-radiometer) data. HDT, 3480/3490, and DCT have proven to be robust and high-performance for their time. As technology advances, as datasets grow, as media ages, and as USGS Digital Library space fills, the USGS must migrate data to newer, more physically compact, and higher performing storage technologies.

1.3 Data reliability

Since the foremost goal of an archive is data preservation, the primary criteria for the selection of the drive technology must be reliability. Several elements contribute to data reliability:

- The number of archival copies: The dependence on the master copy, and the level of risk rise when a working copy is not robust. Any of the reviewed technologies would require a master copy though some would rely on it more. Note that the master and working copies need not be on similar media, though generation and recovery of a working copy is simplified if the storage capacities are similar. All USGS archives must have both working and master copies, and an offsite copy is desirable. Note that a slightly less reliable drive can be used if there are a sufficient number of copies of the archive.

- The storage location and environment: This is a constant for all of the technologies assessed since any would be stored in a secure and climate-controlled environment.

- The composition of the media: Some media compositions last much longer than others.

- Tape handling within the drive: This characteristic defines how a tape is handled by the drive, whether contact is made with the recording surface, how many passes are required to read or write an entire tape, and the complexity of the tape path.

- Error handling: The ideal drive minimizes data loss through CRC (Cyclic Redundancy Check) or other data recovery methods, and allows data to be read after skipping over an error. Though error detection upon write is required, additional attention to data recovery upon read is a higher priority.

- Primary Market: This criterion describes the target market of a drive, and the characteristics of drives within that market. A drive targeted to the backup market would be designed for write many/read rarely, therefore more emphasis is placed on detecting errors upon write. A drive targeted to the archival market would be designed for write once/read few and more emphasis is placed on detecting and correcting errors upon read. A drive targeted to the Enterprise market would be designed for write many/read many and equal emphasis is placed on detecting errors upon read and write. Ideally, all archives would be written to a drive designed for the archive market, but none are currently available. Most vendors would argue that their products are archive devices, but if forced to choose their primary market nobody would choose the very limited archive market. With proper handling and multiple copies, any of the drives evaluated in this report could be deployed as archive drives.

Primary Market	Reliability	Usage	Driving Design Factors
Backup	Moderate	Write many, read rarely	Cost, capacity, speed
Enterprise	High	Write many, read many	Designed for continual use, often with robotics
Archive	High	Write once, read few	Long term reliability

Table 1-2 Tape Drive Markets and Characteristics

The reliability of a long-term archive technology relates primarily to the long-term viability of the recorded media. Since it is wise to implement a technology early enough in its life cycle that drives can be kept viable through the lifetime of a given media (or replaced with newer backward-compatible models), a definitive leader in reliability is difficult to determine except in retrospect. This study bases the reliability assessment on past experience with the vendor and their products, on specifications, on the experiences of others, or experience gained from benchmarking.

Based on a USGS study of DLT 7000 errors, it is suspected that the way that Quantum implemented servo tracks leads to an increase in data loss upon each occurrence, as compared to 3490. When an error occurred, it frequently appeared in several places on the tape (presumably in the same linear location, across multiple tracks) and there was more data loss at a given location as compared to 3490. In many cases, data could not be recovered past the error, as is typically possible with 3490. Lack of servo track redundancy is suspected as the cause.

SDLT (Super Digital Linear Tape) may provide an improvement since the servo tracks are located on the back of the tape, using indelible optical markings. The USGS has not used SDLT extensively enough to determine whether optical servo tracks have improved SDLT reliability over DLT 7000, but they likely have. On the first LTO tape tested at the USGS, a problem occurred which is reminiscent of the DLT 7000 errors.

StorageTek 9940 uses serpentine recording but uses many fewer passes than either LTO or SDLT. In addition, 9940 drives do not touch the recording surface, and redundant servo tracks are provided. Experience with 3480, 3490, 9840, 9940A and 9940B has shown StorageTek products to be very reliable. The StorageTek D3 helical scan drive was problematic and was discontinued quickly. On two occasions, 9840 tapes that encountered unrecoverable errors were sent to StorageTek for recovery. One tape was recovered, but the other was unrecoverable due to cartridge contamination.

1.4 Technologies selected for consideration

The criteria used in determining which technologies should be considered were:

1. The technology must be currently available and shipping in order to be considered in the final analysis. It also must be the latest drive in the line. Other technologies may be mentioned in the study if they meet the other criteria listed here and are projected to become available by the end of calendar 2003.

2. The technology must hold more than 50GB of data.

3. The technology must have a write transfer rate of at least 10 MB/sec. SDLT 220 was allowed into the field with a measured transfer rate less than 10 MB/sec because the advertised rate is 11 MB/sec.

4. The technology must use a media that can remain readable for at least 10 years in a controlled environment. The lifetime of 10 years was selected since it is the longest that a media technology would conceivably be used before space and transfer rate concerns would dictate a move to a new technology. Metal particle and optical media meet this criterion, but chromium dioxide (CRO2) would not since it is good for only 5 years.

5. The technology must not use helical scan technology. This is based on years of bad experience with helical scan 8mm, 4mm, and StorageTek D3 at the USGS. Helical drive reliability and tape wear are a concern due to the constantly moving heads that contact the tape, as well as a complicated tape path. Exabyte drives were so problematic that double the number of drives were needed since half were broken at any given time. The StorageTek D3 drive was discontinued soon after release.

6. The technology must not be hampered by a poor reliability history.

The currently available drive technologies selected for final consideration are:

1. StorageTek 9940B
2. LTO2 (Linear Tape Open)
3. SuperDLT 320

The following technologies are noted in the tables for comparison purposes, but not considered in the final analysis for the reasons noted:

1. LTO1 (Has been replaced by the much more capable LTO2)
2. SuperDLT 220 (Has been replaced by the much more capable SuperDLT 320)
3. SuperDLT 640 (Not yet available, projected specifications from vendor or by estimation)
4. LOTS LaserTAPE (Not yet available, projected specifications from vendor or by estimation)
5. IBM (International Business Machines) 3590H (Cost prohibitive based on scenario costs)

1.5 Dismissed technologies

The following technologies were dismissed from further analysis or consideration for the reasons listed.

1.5.1 CD-ROM, DLT 8000, QIC, and Erasable Optical (EO)

This category includes technologies that are low capacity, low performance, or aged. All of these products have been available for some time, but can immediately be dismissed based on obvious limitations in performance, capacity, or reliability. These products are clearly not a good fit for large digital archives.

1.5.2 Exabyte VXA2 and Mammoth 2

Exabyte has evolved its early helical scan technology into two product lines: VXA2 with a native capacity of 80 GB a native transfer rate of 6 MB/sec and the Mammoth 2 with a native capacity of 60 GB and a native transfer rate of 12 MB/sec. While media costs are low, transfer rates are acceptable, and company stability is moderate, helical scan technology has not proven reliable over time.

1.5.3 Sony Super-AIT

AIT (Advanced Intelligent Tape) is an evolutionary step up from the 8mm helical scan drives made popular by Exabyte. The most recent generation of the AIT yields a native capacity of 500 GB, and a 30 MB/sec transfer rate. Like the current Exabyte offerings, helical scan AIT raises serious reliability concerns.

1.5.4 DVD

DVD (Digital Video Disc) seems promising from the standpoint of longevity of the media. However, low capacity per media, low transfer rates, lack of media protection, no single standard, and high media costs add up to a product that simply would not work for high volume archival use.

1.5.5 Other technologies

Several high capacity optical disk technologies have been in the development phase for the past few years. Of the 100+ GB technology proposals that have appeared in trade journals and at conferences, to date none are shipping products, and several have vanished.

Other high-tech examples of future technologies such as holographic storage or bio-storage will not mature for several years.

2.0 Technical Assessment

2.1 Analysis

StorageTek 9940B:

Advantages:

- USGS experience with 9940A and 9940B drives at EDC has shown them to be more reliable than DLT. Past USGS experience with StorageTek 3480/3490 compatible drives has shown StorageTek products to be very reliable. Advantages of the 9940B include 'wider' tracks (16 tracks per pass instead of 8) to reduce serpentine passes, and air bearings that allow the tape to float past the head without contact.

- 9940 is targeted to the Enterprise Storage market where data viability, speed, and capacity are more important than cost.

- 9940 was designed as a robust storage media, with the tape cartridge and drive built to withstand constant and/or frequent use in a robotic environment. The 9940 drives are compatible with the USGS StorageTek silos and excel in a robotic environment due to their durability.

- The USGS offline Digital Library shelving and tape carriers used for 3480/3490 work with 9940.

Disadvantages:

- StorageTek is the sole manufacturer of 9940B.

- StorageTek recently indicated that the 9940B is the last of the 9940 series, as they have reached the limits of metal particle technology in the 9940 design. They are working on a new product (9950?) with the first new drive at 500 GB native capacity shipping in 2004 or 2005. The drive will use new media and it is unclear whether this drive would be backward compatible.

- StorageTek only sells the non-Silo 19" rack-mount version of the 9940B in pairs.

- The drives are relatively expensive.

Notes:

- The usable capacity may vary between cartridges. The USGS attained a capacity of 193.03 GB per tape. CERN (Conseil European pour la Recherché Nucleaire) was able to write 208 GB per cartridge on all 10 tapes they tested (http://cscct.home.cern.ch/cscct/T9940B.ppt). StorageTek indicates that capacity may vary by 10% between batches of tapes and that the tapes likely came from the same batch. It is unclear whether 10% meant plus or minus 10% (180 to 220GB) or a 10% window (190 to 210GB).

- While the projected follow-on 9950 will take different media, it is anticipated that StorageTek will continue the tradition of using the same physical cartridge dimension so that existing robotic libraries can accept the new media without modification. This should also ensure compatibility with offline shelving.

LTO1:

Advantages:

- LTO1 has enjoyed phenomenal growth from the day of release, continuing through the recent slowdown in the IT sector despite contraction of the tape industry as a whole.

- LTO has a 67% market share, with 250,000 drives installed worldwide, compared to 150,000 SDLT drives.

Disadvantages:

- LTO1 is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write-many/read-rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.

- LTO1 may suffer from the same data loss characteristics as the DLT 7000. Reliability is a concern since one end-to-end read/write would incur 48 passes.

- LTO1 was co-developed by Seagate, IBM, and HP (Hewlett Packard). This type of deployment makes it possible for each vendor to interpret the specifications differently, and to design drives which may have incompatibilities. Though they may test interoperability, competition encourages differentiation. There have been hints of cross-brand problems mentioned on the Internet, and by one reseller. Because of this concern, if LTO were selected it would be advisable to utilize only one vendor.

- LTO was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use. Although STK recently added the capability for their large silo to handle both LTO and 9940/9840, the robotic arm had to be slowed down since the thin shell of the LTO could not take the grip pressure necessary to keep cartridges from flying out of the gripper when the arm is at full speed.

- The first tape written on LTO at the USGS encountered an unrecoverable read error and data past the point of the error was unrecoverable. The LACS (Landsat Archive Conversion System) project experienced a tape write error on a tape that had been written 50 times, and read 25 times - which is well under the tape usage cycles specified by the manufacturer.

Notes:

- There is very little USGS experience with LTO. The USGS procured an LTO drive that is being used for LACS testing. During the initial tests, the LTO performed very close to the specified speed and capacity.

- Repair would require a return to the vendor service center. Due the typical downtime associated with this method of service, spare drives would be required.

LTO2: (in addition to LTO1 notes above)

Advantages:

- After market saturation, LTO2 is projected to have lowest media cost per TB.
- Backward read/write compatible with LTO1. This means that the LTO2 drive can read and write LTO1 cartridges in the LTO1 density. All future drives are slated to be able to read any previous generation of tape.

Disadvantages:

- LTO2 uses a different media than LTO1, which will drive media costs up until market saturation brings costs back down.
- Designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use.

Notes:

- IBM and HP are shipping LTO2 now, although IBM is shipping engineering units only to partners. The price quoted in chapter 4 is for an HP drive.
- The third and fourth generations of LTO Ultrium have been projected but not scheduled. LTO3 and LTO4 will have native capacities of 400 and 800 GB and native transfer rates of 80 and 160 MB/sec. The LTO consortium does not estimate dates for future products.

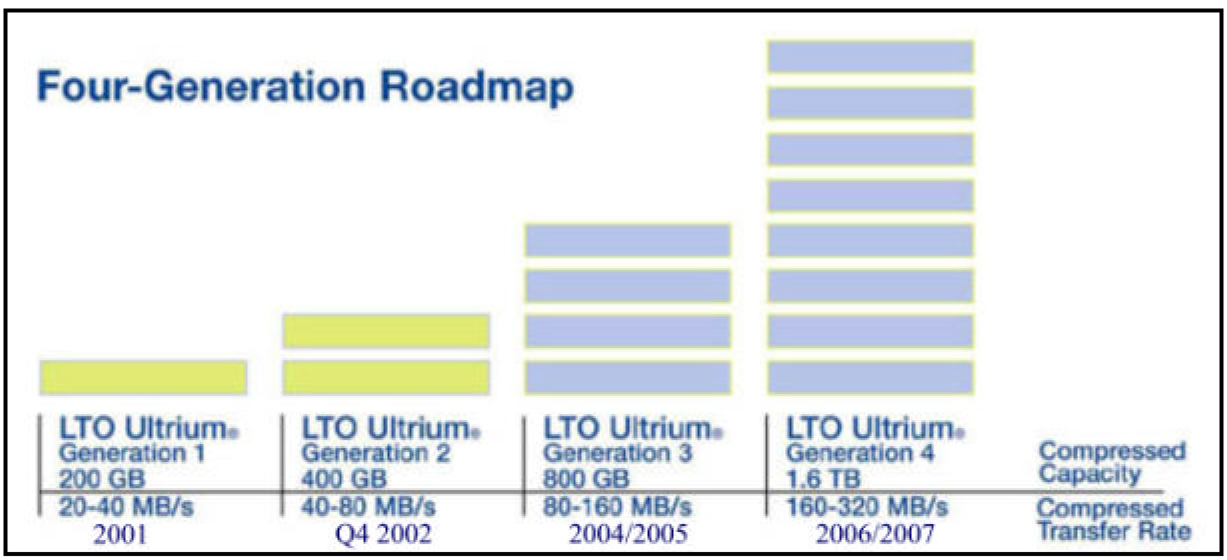

Figure 2-1 LTO Roadmap (with 2:1 compression)

SDLT 220:

Advantages:

- SDLT has enjoyed very wide market saturation due to its history in the backup market, although LTO has overtaken SDLT.

Disadvantages:

- SDLT is targeted to the backup market where speed, capacity, and cost are more important than long-term viability of the data. Since backup tapes are write-many/read-rarely, errors would likely show up in a write pass where they can be worked around (rewrites) or the media discarded.
- There is little USGS experience with SDLT, although the LP-DAAC has archived MODIS data on SDLT 220. During initial tests, SDLT performed poorly – not coming close to the specified speed and capacity.
- Tape wear is a concern since one end-to-end read/write incurs 56 passes over the heads.
- Media costs for SDLT cartridges will probably remain higher than the competition due to the optical servo track imprinted on the back of the tape.
- SDLT was designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use.
- Quantum licensed drive manufacturing to other vendors, then later bought out those companies.

Notes:

- Repair would require a return to the vendor service center. Due the typical downtime associated with this method of service, spare drives would be required.

SDLT 320: (in addition to SDLT 220 notes above)

Advantages:

- SDLT 320 drives are priced lower than LTO2.
- SDLT 320 uses the same media as SDLT 220, which will likely mean a continuing slow decrease in media cost.

Disadvantages:

- Transfer rates have improved over the SDLT 220, but are still lower than LTO2.
- Capacity is significantly lower than LTO2, its primary competitor.
- Designed as a moderate usage storage media, with the tape cartridge and drive not built to withstand constant or frequent use.

Notes:

- The product roadmap below calls for a 320 GB drive at 32MB/sec by late 2003, 600 GB at 64 MB/sec by the second half of 2004, and 1.2 TB (Terabytes) at 100 MB/sec in 2006.

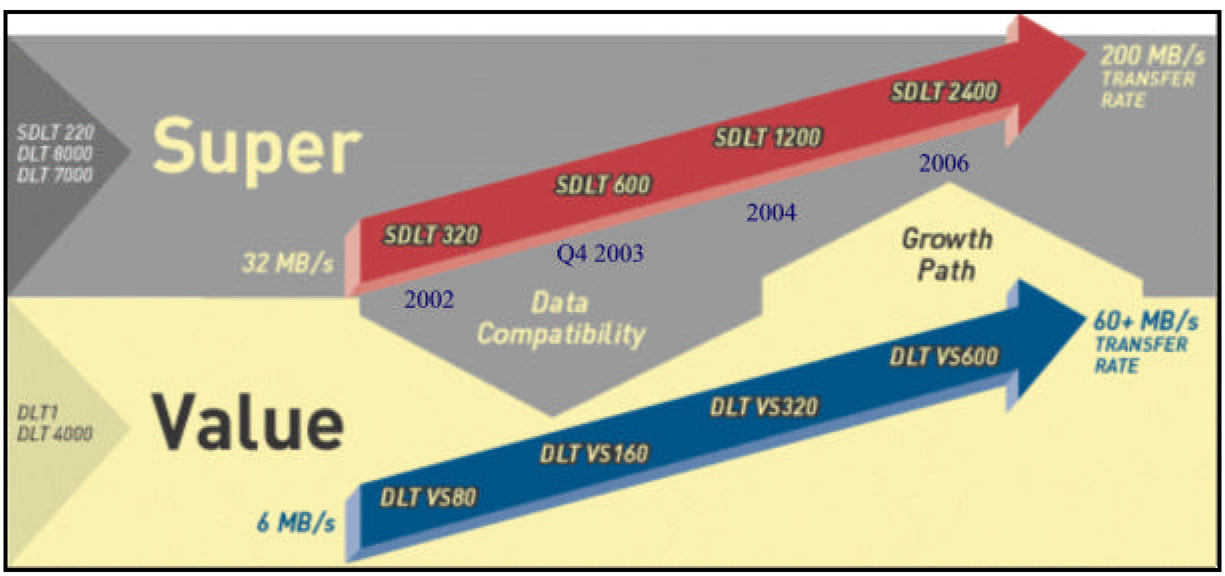

Figure 2-2 SuperDLT Roadmap (with 2:1 compression)

SDLT 640: (in addition to SDLT 220 and 320 notes above)

Advantages:

- Capacity will exceed 9940B and LTO2.

Disadvantages:

- SDLT 640 is slated to use different media than SDLT 220 and 320, virtually ensuring that media costs will be significantly higher until market saturation drives the price down.

- Transfer rates will be lower than 9940B and LTO2.

Notes:

- Due for release in mid 2003.

- Capacity has been stated as 300 GB and 320 GB, though the 300 GB figure comes from the latest roadmap.

IBM 3590H:

Advantages:

- Based on the very reliable 3480 and 3490.
- Good transfer rate.
- IBM plans to keep the same media through the next 4 generations.
- 3590 was designed as a robust storage media, with the tape cartridge and drive built to withstand constant and/or frequent use.

Disadvantages:

- Tape capacity is low compared to the other technologies reviewed. Capacity is based on extended length tapes.

Notes:

- IBM seems to be one generation behind the competition on capacity.
- The exact tape drive model number is 3590-H11.
- IBM has a product roadmap showing future 3590 drives having capacities of 200 GB, 400 GB, and 1000 GB native, and transfer rates reaching 160 MB/sec. No timeline is provided but a prototype of the 1000 GB was running as of May 2002.

LOTS LaserTAPE:

Advantages:

- Capacity is slated to be 1 TB in a 3480/3490/9940/3590 sized cartridge.
- Net transfer rate is stated to exceed 40 MB/sec.
- WORM (Write Once, Read Many) optical technology with 100-year longevity projected.

Disadvantages:

- Projected to be expensive
- Will initially only work in a Grau robot or stand-alone.
- LaserTAPE has been pre-announced for over 10 years. In June 2002, they claimed a ship date in December 2002. Because of the stop-and-go design and marketing, they will struggle to earn credibility. As of January 17th, 2003 the Lots website is no longer operational which may indicate that they did not receive financing necessary to continue operation.

Notes:

- The LaserTAPE was slated for release December 2002. On 12/13/02, LOTS stated that the Alpha drive testing was successful and they were now seeking $5M to build beta drives. If they can obtain funding by the end of January 2003, they intend to ship beta drives in the summer and production drives in the fall. They stated that the alpha drive exceeds the specified 40 MB/sec transfer rate and 1 TB capacity.
- With an advertised media life of 100+ years, one must question how tapes would be read in 25 years, 40 years, and beyond. How likely will it be for the company to succeed, last 40 years, and still have backward read capability? Space and performance issues would likely drive migration before the media end-of-life, probably in no more than 25 years but it would still be beneficial to have confidence in the data even after 25 years.

3.0 Tables

3.1 Design criteria

The design criteria and target market of a drive are interrelated. Tape drives such as LTO and DLT are targeted to the backup market as demonstrated by their marketing. The 9940B and 3590H are targeted to the Enterprise (data center) market. The LaserTAPE is the only drive in this study that is targeted to the Archive market.

A drive targeted to the backup market is designed for write many/read rarely and more emphasis is placed on detecting errors upon write. Backup drives are typically built for speed, capacity, and low cost. A drive targeted to the Enterprise market is designed for write many/read many use in a robotic library or auto-stacker and equal emphasis is placed on detecting errors upon read and write. A drive targeted to the archival market is designed for write once/read few and more emphasis is placed on detecting and correcting errors upon read.

Enterprise and Archival drives are typically built for reliability, with speed and cost a secondary factor. All drives attempt error detection and recovery upon both read and write, but an archival drive design typically places more importance on read data recovery since data may no longer be available. Conversely, a backup drive places more importance on write error detection since the data is still available and can be easily rewritten.

Note that rankings could not be determined where information was unavailable. The formula used to rank Design Criteria was:

((100-passes)/10)+
(error factor/2)+
(construction 3=moderate, 5=high usage)+
(head contact 3=yes, 5=no)
/ 2.54 (to adjust the highest rank to 10)

Technology	Serpentine tracks/ Passes	Target Market	Tape Composition	Uncorrected Error Rate	Cartridge Construction Rating	Head Contact	Ranking
9940B	576/36	Enterprise	Advanced MP	1x10^{18}	High usage	No contact	10.0
LTO 1	384/48	Backup	Metal Particle	1x10^{17}	Moderate usage	Contact	7.7
SDLT 220	448/56	Backup	Advanced MP	1x10^{17}	Moderate usage	Contact	7.4
SDLT 320	448/56	Backup	Advanced MP	1x10^{17}	Moderate usage	Contact	7.4
LTO 2	512/64	Backup	Metal Particle	1x10^{17}	Moderate usage	Contact	7.1
SDLT 640	Unknown	Backup	Advanced MP	Unknown	Moderate usage	Contact	
3590H	384/24	Enterprise	Advanced MP	Unknown	High usage	Contact	
LaserTAPE	Unknown	Archival	Optical	1x10^{17}	Unknown	No contact	

Table 3-1 Design criteria and target market
(Green indicates drives considered)

3.2 Transfer Rate

Transfer rate is important since it dictates how many months or years will be required to migrate an archive, and how fast a production system may generate products from the archive media. Since the typical read transfer rate of the DCT drive is 10 MB/sec, the requirement is to match this rate. Although 10 MB/sec is the minimum, it is desired to double this transfer rate to 20 MB/sec. The SDLT 220 was included in the study since the advertised rate is above the minimum. Since much of the data archived at the USGS is not compressible, all transfer rates are native (uncompressed).

The ranking was determined by adding the actual/estimated read and write rates for each drive, setting the ranking for the fastest drive to 10, then ranking the others against the leader. As an example, a drive having half of the total read/write transfer rate of the leader would be ranked 5.

Tape Drive Technology	Advertised Native Rate	Actual/estimated Native Write Transfer Rate	% Of Adv.	Actual/estimated Native Read Transfer Rate	% Of Adv.	Ranking
9940B	30 MB/sec	25.46 MB/sec	84.8%	28.55 MB/sec	95.2%	10.0
LTO 2	30 MB/sec	23.83 MB/sec	79.4%	18.52 MB/sec	61.7%	7.8
SDLT 320	16 MB/sec	14.79 MB/sec	92.4%	16.09 MB/sec	100%	5.7
LTO 1	16 MB/sec	14.66 MB/sec	91.6%	10.32 MB/sec	64.5%	4.6
SDLT 220	11 MB/sec	8.12 MB/sec	73.8%	6.35 MB/sec	57.7%	2.7
SDLT 640	32 MB/sec	23.60 MB/sec est.	73.8%	18.46 MB/sec est.	57.7%	
3590H	14 MB/sec	14.20 MB/sec	100%+	14.00 MB/sec est.	<NA>	
LaserTAPE	40 MB/sec	40 MB/sec claimed	100%	40 MB/sec claimed	100%	

Table 3-2 Transfer rates
(Green indicates drives considered)

3.3 Capacity

A secondary requirement is to conserve archive rack space by increasing per media capacity. The current archive media of choice at the USGS is DLT 7000 at 32 GB per tape. The new minimum requirement is 50 GB, with 100 GB or more desired. All of the reviewed technologies meet the 50 GB requirement. Since much of the data archived is not compressible, all capacities are native (uncompressed).

The ratings were determined by calculating each as the percentage of the highest capacity drive, on a scale of 1 to 10, with the highest capacity as a 10.

Tape Drive Technology	Advertised Native Capacity	Measured/Estimated Native Capacity	% Of Advertised Capacity	Ranking
LTO 2	200 GB	197.00 GB	98.5%	10.0
9940B	200 GB	193.03 GB	96.5%	9.8
SDLT 320	160 GB	153.00 GB	95.6%	7.8
SDLT 220	110 GB	98.83 GB	89.8%	5.0
LTO 1	100 GB	97.75 GB	97.7%	5.0
SDLT 640	300 GB	269.40 GB estimated	89.8%	
3590H	60 GB			
LaserTAPE	1000 GB	1000 GB claimed	100% claimed	

Table 3-3 Storage Capacities
(Green indicates drives considered)

3.4 Cost Analysis

Table 3-4 shows the relative drive and media costs, maintenance costs, and the cost per Terabyte for media. Note that the price of LTO2 media and drives is expected to drop significantly within six months, so a projection of those lower costs has been included. Rankings were established by setting the cheapest (drive, maintenance, media) to 10 then rating each of the others against the lowest cost.

Technology	Drive $/each	Annual Maintenance	Media $/each	Media $/TB	Ranking Drive Cost	Ranking Maint Cost	Ranking Media Cost
9940B	$30,000	$4,500	$78	$403	1.0	1.0	10.0
LTO 2	$5,549	$832	$119	$604	5.0	5.0	6.6
SDLT 320	$3,700	$555	$95	$621	7.6	7.6	6.5
LTO 1	$3,568	$535	$63	$644	7.8	7.8	6.3
SDLT 220	$2,800	$420	$95	$961	10.0	10.0	4.2
SDLT 640	$6,000 est.	$900	$120 est.	$445			
3590H	$33,945	$5,091	$102	$1,700			
LaserTAPE	$27,000 est.	$4,050	$200 est.	$200			

Table 3-4 Drive, maintenance and media costs
(Green indicates drives considered)

3.5 Scenarios

Table 3-5 shows the total drive and media cost for four scenarios. Table 3-6 shows the estimated migration times for each scenario, which is the time to write the data once and to read the data once as a verification step, using all drives. Note that a minimum of two drives are required for redundancy. These scenarios presume that each dataset or project stands on their own, but pooling resources for multiple datasets can mitigate cost. Note that prices are expected to drop considerably within six months after product introduction. Rankings were established by setting the cheapest or fastest to 10 then rating each of the others against the lowest cost or lowest completion time.

Technology	10 TB 2 drives	100 TB 4 drives	250 TB 6 drives	500 TB 8 drives	Ranking
SDLT 320	$13,610	$76.900	$177,450	$340,100	10.0
LTO 1	$14,646	$80,812	$185,618	$354,824	9.6
LTO 2	$17,138	$82,596	$184,294	$346,392	9.6
SDLT 220	$16,050	$108,980	$259,570	$506,260	6.8
9940B	$73,030	$178,300	$307,750	$477,500	5.8
SDLT 640	$18,250	$72,100	$152,650	$277,700	
3590H	$95,072	$326,144	$659,216	$1,162,288	
LaserTAPE	$64,100	$144,200	$236,300	$348,400	

Table 3-5 Scenario costs (drives, 1 year maintenance, media)
(Green indicates drives considered)
(Yellow indicate s scenario selected for criteria table)

Technology	10 TB 2 drives	100 TB 4 drives	250 TB 6 drives	500 TB 8 drives	Ranking
9940B	103	516	860	1290	10.0
LTO 2	133	666	1111	1666	7.7
SDLT 320	180	901	1502	2253	5.7
LTO 1	229	1147	1911	2867	4.5
SDLT 220	390	1949	3248	4872	2.6
SDLT 640	134	670	1117	1676	
3590H	196	978	1630	2445	
LaserTAPE	69	347	579	868	

Table 3-6 Migration Times in hours (1 write, 1 read)
(Green indicates drives considered)
(Yellow indicates scenario selected for criteria table)

3.6 Vendor analyses

Table 3-7 is intended to provide an analysis of each company and the stability of each technology. Since the final criteria table lists only technologies (but not vendors), the points for the three LTO vendors were averaged. The Producer Ranking is based on the number of manufacturers, with the leader adjusted to 10 by multiplying by 3.33. The longevity rankings were determined by the following formula:

(company age/10) +
(technology age))
 * 1.25 (to adjust the highest rank to 10)

Company	Technology	Number Of Manufacturers	Years in business	Technology age in years	Producer Ranking	Longevity Ranking
LTO consortium	LTO	3	60 avg	2	10.0	10.0
Quantum	SDLT	1	23	2	3.3	5.4
StorageTek	9940	1	34	2	3.3	6.7

Table 3-7 Vendor Analyses

3.7 Drive compatibility

Table 3-8 shows the level of inter-generation drive compatibility as well as the future drives planned. The column "Previous Generations Read" indicates how many previous generations are read by the generation indicated. The column "Future Generations Planned" indicates the number of generations planned in the current drive family, following the drive indicated. The column "Future Generations Compatible" indicates how many future generations are announced to be compatible with the drive indicated. The ranking was determined by the following formula:

(Previous Generations Read + Future Generations Planned + Future Generations Compatible)
 * 1.25 (to adjust the highest rank to 10)

Technology	Previous Generations Read	Future Generations Planned	Future Generations Compatible	Ranking
SDLT 320	2	3	3	10.0
SDLT 220	1	4	3	10.0
LTO 1	0	3	3	7.5
LTO 2	1	2	2	6.2
9940B	1	0	0	1.2

Table 3-8 Drive Compatibility
(Green indicates drives considered)

3.8 Ranking summary

The ranking summary provides a quick reference to the rankings. X

Drive	Design Criteria	Transfer Rate	Capacity	Drive/ Maint Cost	Media Cost	Scenario Cost	Scenario Time	Producer	Vendor Analyses	Drive Compat.
9940B	10.0	10.0	9.8	1.0	10.0	5.8	10.0	3.3	6.7	1.2
LTO1	7.7	4.6	5.0	7.8	6.3	9.6	4.5	10.0	10.0	7.5
LTO2	7.1	7.8	10.0	5.0	6.6	9.6	7.7	10.0	10.0	6.2
SDLT 220	7.4	2.7	5.0	10.0	4.2	6.8	2.6	3.3	5.4	10.0
SDLT 320	7.4	5.7	7.8	7.6	6.5	10.0	5.7	3.3	5.4	10.0

Table 3-9 Ranking Summaries
(Blue indicates the highest ranking of the drives that made the cut)

4.0 Conclusions and Recommendations for USGS Offline Archiving Requirements

4.1 Weighted Decision Matrix

The following table provides a weighted analysis of the three drives considered. The criteria emphasize the importance of traits contributing to data preservation. The sponsor made the final decision regarding which criteria to use and the relative weighting of the criteria. The columns in green are relative ratings for each technology. The columns in yellow are calculated by multiplying the relative weight by the relative rating. The following describe each criterion:

- Design Criteria (Reliability of media): This criterion describes the ability of the media to remain readable over time. Included in this criterion is the number of passes per full-tape read or write, cartridge construction, uncorrected bit error rate (BER) and amount of head contact. (See table 3-1)
- Capacity: This criterion describes the measured or estimated capacity per cartridge, which is typically less than the advertised capacity. (See table 3-3)
- Media cost/TB: This criterion is a rating of the relative cost per Terabyte for media using the measured or estimated capacity rather than advertised capacity. (See table 3-4)
- Maintenance cost: This criterion is the relative rating of drive maintenance cost. Note that at this point, this rating is the same as the drive cost rating since maintenance is estimated at 15% of the drive cost annually. (See table 3-4)
- Compatibility: This criterion describes the likelihood that the drive technology will continue to evolve and the extent to which future drives will have backward read capability. This will give an indication of the ability to maintain drives that can read an aging archive. (See table 3-8)
- Transfer rate: This criterion describes the aggregate read and write transfer rate, which is typically less than the advertised transfer rate. (See table 3-2)
- Drive cost: This criterion is the rating of relative cost of each drive at the lowest currently available price. (See table 3-4)
- Vendor analyses: This criterion is the rating of the viability of the vendor and technology. (See table 3-7)
- Single vendor: This criterion is the rating of the vendor single-point-of-failures risk. (See table 3-7)
- Scenario cost: This criterion is the rating of the cost of scenario #3. This includes media cost/TB, drive cost, and maintenance cost. The measured or estimated capacity is used rather than advertised capacity. This criterion is rated on a scale of 1 to 10 with the most expensive being a 1 and the least expensive a 10. (See table 3-5)
- Scenario time: This criterion is the rating of the time to complete a migration write and verify operation using the measured or estimated transfer rate and capacity rather than advertised figures. This criterion is rated on a scale of 1 to 10 with the shortest time rated a 10 and the longest time rated a 1. (See table 3-6)

Note that in the decision matrix spreadsheet below, not all criteria have been selected for the final analysis of this trade study. These unused criteria were left in the spreadsheet so that others may insert the criteria weights for their specific application.

Selecton Criteria	Relative weight	9940B	LTO2	SDLT 320	9940B	LTO2	SDLT 320
Design criteria (reliability)	33	10.0	7.1	7.4	330.0	234.3	244.2
Capacity	10	9.8	10.0	7.8	98.0	100.0	78.0
Media cost/TB		10.0	6.6	6.5	0.0	0.0	0.0
Maintenance cost		1.0	5.0	7.6	0.0	0.0	0.0
Compatibility	15	1.2	6.2	10.0	18.0	93.0	150.0
Transfer rate		10.0	7.8	5.7	0.0	0.0	0.0
Drive cost		1.0	5.0	7.6	0.0	0.0	0.0
Vendor analyses		6.7	10.0	5.4	0.0	0.0	0.0
Single vendor (not)	10	3.3	10.0	3.3	33.0	100.0	33.0
Scenario cost	17	5.8	9.6	10.0	98.6	163.2	170.0
Scenario time	12	10.0	7.7	5.7	120.0	92.4	68.4
Total Weighted Score					**697.6**	**782.9**	**743.6**

Table 5-1 Decision matrix

4.2 Conclusions and notes

- The clear point leader is LTO2.
- The primary criterion is reliability and 9940B leads in this crucial category. These findings do not mean that the SDLT 320 and LTO2 are unreliable, just less reliable and less durable than the 9940.
- As any drive saturates the market, media costs drop. It is projected that by mid to late 2003, LTO2 media costs will drop considerably, and there may also be a significant drop in drives cost.

4.3 Recommendations

1. It is advised that the USGS move quickly to procure an HP or IBM Ultrium 460 (LTO2) drive for prototyping.
2. Presuming the results of these tests prove the LTO2 technology to be capable, it is advised that the LTO2 technology be adopted for the next generation offline archive media technology.
3. In order to reduce risk, the USGS should develop a strategy for storing datasets on multiple technologies. An example of this would be to store a working copy of the Landsat dataset on nearline 9940B, and offline master and offsite copies on LTO2. The teaming of 9940B and LTO2 would provide ease of recovery since their capacities are nearly identical and could easily be written with identical contents on each volume. This teaming strategy partially mitigates the risks of one or the other company failing or the technology being retired prematurely.

4. Since this study does not directly assess nearline (robotic) technologies, it is advised that the USGS continue plans to deploy 9940B and its successors for robotic storage of archive working copies.

5. The USGS is advised to continue to deploy RAID technology for smaller datasets when possible and feasible. RAID disk will continue to be an option for small to moderate dataset working copies, as long as there is an offline master copy. As disk prices decrease, it will become more feasible to store working copies for large frequently-accessed datasets on RAID.

6. The USGS should continue to monitor LOTS LaserTAPE, and the optical market in general as optical has a much longer shelf life.

7. The USGS should continue to monitor IBM Magstor 3590 developments. As the tape density increases, costs per Terabyte will decrease.

8. The USGS should continue to monitor SuperDLT developments. In order to stay viable, SuperDLT costs and features must be competitive with LTO.

9. The USGS should plan to update this trade study at least annually.

Abbreviations and Acronyms

AIT	Advanced Intelligent Tape
CD-ROM	Compact Disc - Read Only Memory
CERN	Conseil European pour la Recherché Nucleaire (European Laboratory for Particle Physics; Geneva, Switzerland)
CRC	Cyclic Redundancy Check
DCT	Digital Cassette tape
DLT	Digital Linear Tape
DVD	Digital Video Disc
EDC	EROS Data Center
EROS	Earth Resources Observation Systems
FYyy	Fiscal Year yy
GB	Gigabytes
HDT	High Density Tape
HP	Hewlett Packard
IBM	International Business Machines
IRIG	InteRange Instrumentation Group (timecode format)
LACS	Landsat Archive Conversion System
LP-DAAC	Land Processes Distributed Active Archive Center
LTO	Linear Tape Open
MB	Megabyte
MODIS	Moderate Resolution Imaging Spectroradiometer
MSS	Multi-spectral Scanner
QIC	Quarter-inch Cartridge
SAIC	Science Applications International Corporation
SDLT	Super Digital Linear Tape
TB	Terabytes
WORM	Write Once, Read Many
TM	Thematic Mapper
TMACS	TMMSS Archive Conversion System
USGS	United States Geological Survey